

Managing AWS PBS Jobs Launched From a PFE

This article provides information on managing your Amazon Web Services (AWS) PBS jobs submitted from a Pleiades front end (PFE).

The PBS server at NAS for managing AWS cloud job is clpbs-01. It coordinates with another PBS server at AWS to get your job burst into and running on AWS, and sending the PBS output/error file back to Pleiades.

Note: The `aws_pbs_host`, `aws_qstat`, and `aws_qdel` scripts described in this article are located under the directory `/u/scicon/tools/bin` on Pleiades. The instructions below assume that you have included `/u/scicon/tools/bin` in your `$PATH`.

Finding Available Queues

To find available queues, type the command:

```
pfe% qstat -q @clpbs-01
```

Among the queues listed in the output, you only have direct access to the *cloud* queue. The other queues, such as *frontend* and *s3op*, are used by NAS-controlled scripts (*aws_fe* and *nas_s3_xxx*) for managing the [AWS dynamic front end](#) and [data at AWS S3](#) on your behalf.

The limits of each queue are configured in the PBS server at AWS. To see them, use:

```
pfe% aws_qstat -Q @`aws_pbs_host`  
pfe% aws_qstat -Q @`aws_pbs_host` --group gid` (to check for a non-default GID)
```

Note: If you get the following message, your PBS server at AWS has likely gone to sleep, for cost saving purposes. Submitting a job or two should wake up your server.

```
Connection timed out  
qstat: cannot connect to server your_aws_pbs_server.nas.nasa.gov (errno=110)
```

Submitting Jobs (qsub)

You must submit jobs to AWS from your Pleiades `/nobackup` filesystem.

From a PFE, you can submit regular batch jobs to the *cloud* queue, but not interactive jobs, using the following command:

```
pfe% qsub -q cloud@clpbs-01 job_script
```

To submit jobs to a non-default AWS GID, add the `-W group_list=gid` option, where `gid` is in the form of `zxxxx` (i.e., the letter "z" followed by four numbered digits):

```
pfe% qsub -q cloud@clpbs-01 -W group_list=gid job_script
```

or add it inside your PBS script:

```
#PBS -W group_list=gid
```

Note: If you get the following message within a few minutes of submitting your job, your PBS server at AWS likely went to sleep, for cost saving purposes; however, by this time it should have been awakened and subsequent jobs should go through.

```
PBS Job Id: xxxx.clpbs-01.nas.nasa.gov
Job Name:   test.pbs
Aborted by PBS Server
Job rejected by all possible destinations
```

Checking Jobs (qstat or aws_qstat)

When a job is submitted from a PFE, it is under the control of the PBS server clpbs-01 at the beginning, and running **qstat** at that time on a PFE will show the job.

```
pfe% qstat -a @clpbs-01
```

The status of a job is H (held) while any files are copied to temporary storage on the AWS cloud. When the copy is completed, the status changes to M, indicating that it was moved to the PBS server at AWS. The job then disappears from clpbs-01. However, you can see past jobs using the following:

```
pfe% qstat -x @clpbs-01
```

After clpbs-01 hands over the control for the job to AWS, the job is visible from your AWS PBS server, using:

```
pfe% aws_qstat @`aws_pbs_host`
pfe% aws_qstat @`aws_pbs_host` --group gid` (to check for a non-default GID)
```

You can get more detailed information about a job using the following:

```
pfe% aws_qstat -f jobid.clpbs-01.nas.nasa.gov@`aws_pbs_host` --group gid`
```

When a job finishes, it may take a few minutes for the PBS output file to show up in your Pleiades directory. You can check whether your job has completed using:

```
pfe% aws_qstat -xqst @`aws_pbs_host` --group gid`
```

Deleting Jobs (aws_qdel)

To delete jobs submitted to AWS from a PFE, use the script **aws_qdel**. Note that you cannot delete these jobs using the standard **qdel** command, because it involves the coordination of two PBS servers.

```
pfe% aws_qdel jobid
or
pfe% aws_qdel jobid.clpbs-01.nas.nasa.gov
```

To delete a job for a non-default AWS GID, use:

```
pfe% aws_qdel --group gid jobid
or
pfe% aws_qdel --group gid jobid.clpbs-01.nas.nasa.gov
```

Location of Your PBS Output/Error Files

Within a few minutes after your job has completed or been terminated, the PBS output/error files are sent back to your \$PBS_O_WORKDIR (where you submitted the job from) on Pleiades.

AWS Job Accounting

AWS charges for the compute instances (CPU/GPU), filesystems, storage, and network for transferring data out of AWS. NAS charges an overhead cost, which is a percentage of the total AWS charge. The Job Costs entry includes both the AWS charge and the NAS overhead charge.

Example Job Accounting Summary

Job Resource Usage Summary for	3552.clpbs-01.nas.nasa.gov
Total Runtime	: 00:03:03
Job Stage In Time (free)	: 00:00:00
Job Startup Time	: 00:01:03
Time Spent In PBS Script	: 00:02:00
Job Stage Out Time	: 00:00:00
Walltime Requested	: 02:00:00
Execution Queue	: AWS Cloud
Charged To	: cstaff
Job Finished	: Wed Jul 3 10:38:09 2019
Instance Types (ondemand)	: 1 m5.xlarge
EBS Usage	: 8577331200 bytes
S3 Usage	: 0 bytes
Charged Bandwidth Usage	: 0 bytes
NAS overhead charge	: 0.000 percent
Job Costs	: \$0.00981639861027

Article ID: 585

Last updated: 04 Dec, 2019

Revision: 19

Cloud Computing -> AWS Cloud -> Managing AWS PBS Jobs Launched From a PFE

<https://www.nas.nasa.gov/hecc/support/kb/entry/585/>